

Search Forms
Search Results

Help**User Searches**

US Pre-Grant Publication Full-Text Database

Preferences

US Patents Full-Text Database

Logout

US OCR Full-Text Database

Database:

EBO Abstracts Database

JPO Abstracts Database

Derwent World Patents Index

IBM Technical Disclosure Bulletins

Freeform Search

Term:

[]
Display: Documents in **Display Format:** Starting with Number **Generate:** Hit List Hit Count Side by Side Image

Search History

DATE: Saturday, December 18, 2004 [Printable Copy](#) [Create Case](#)

Set Name **Query**

side by side

Hit Count **Set Name**

result set

DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR

<u>L15</u>	L14 and attribute with valu\$ same string	25	<u>L15</u>
<u>L14</u>	L13 and query	799	<u>L14</u>
<u>L13</u>	("on-line analytical mining" or "online analytical mining" or "olap")	1161	<u>L13</u>
<u>L12</u>	"on-line analytical mining" or "online analytical mining" or "olap"	1161	<u>L12</u>
<u>L11</u>	"on-line analytical mining" or "online analytical mining"	3	<u>L11</u>
<u>L10</u>	705.clas.	30759	<u>L10</u>
<u>L9</u>	705/5	888	<u>L9</u>
<u>L8</u>	382/229	899	<u>L8</u>
<u>L7</u>	382.clas.	44553	<u>L7</u>
<u>L6</u>	707.clas.	24011	<u>L6</u>
<u>L5</u>	707/5	3383	<u>L5</u>
<u>L4</u>	707/4	4110	<u>L4</u>
<u>L3</u>	707/3	7511	<u>L3</u>
<u>L2</u>	707/2	4404	<u>L2</u>
<u>L1</u>	707/1	7261	<u>L1</u>

END OF SEARCH HISTORY

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)

End of Result Set

 [Generate Collection](#) [Print](#)

L11: Entry 3 of 3

File: USPT

May 13, 2003

US-PAT-NO: 6564197

DOCUMENT-IDENTIFIER: US 6564197 B2

TITLE: Method and apparatus for scalable probabilistic clustering using decision trees

DATE-ISSUED: May 13, 2003

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Sahami; Mehran	Mountain View	CA		
John; George Harrison	San Mateo	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
E.piphany, Inc.	San Mateo	CA			02

APPL-NO: 09/ 304509 [\[PALM\]](#)

DATE FILED: May 3, 1999

INT-CL: [07] [G06 N 5/02](#)

US-CL-ISSUED: 706/55; 707/6, 703/2

US-CL-CURRENT: 706/55; 703/2, 707/6

FIELD-OF-SEARCH: 706/55, 345/440, 382/224, 703/2, 707/6

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<input type="checkbox"/> 5787274	July 1998	Agrawal et al.	395/613
<input type="checkbox"/> 5799300	August 1998	Agrawal et al.	707/5
<input type="checkbox"/> 5809499	September 1998	Wong et al.	707/6
<input type="checkbox"/> 5930392	July 1999	Ho	382/224
<input type="checkbox"/> 6128587	October 2000	Sjolander	703/2
<input type="checkbox"/> 6233575	May 2001	Agrawal et al.	707/6

6278464

August 2001

Kohavi et al.

345/440

FOREIGN PATENT DOCUMENTS

FOREIGN-PAT-NO

0 067 057

PUBN-DATE

December 1982

COUNTRY

EP

US-CL

OTHER PUBLICATIONS

- Frakes, W.B. et al. 1992. Information Retrieval, Data Structures & Algorithms, pp. 419-442.
- Fayyad, U.M. et al. 1996. Advances in Knowledge Discovery and Data Mining, pp. 152-181.
- Gray, J. et al. 1996. Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals. pp. 29-53.
- Chow, C.K. 1968. IEEE Transactions On Information Theory, vol. IT-14, No. 3, Approximating Discrete Probability Distributions With Dependence Trees. pp. 462-467.
- Fisher, D.H. 1986. Unsupervised Concept Learning and Discovery, Knowledge Acquisition Via Incremental Conceptual Clustering. pp. 267-283.
- Chickering, D.M. 1996. Learning Bayesian Networks is NP-Complete. pp. 121-130.
- John, G.H., Lent, B. 1997. American Association For Artificial Intelligence, SIPPING From the Data Firehose. pp. 199-201.
- Sahami M. 1999. Using Machine Learning To Improve Information Access, Dissertation, Stanford University Dec. 1998.
- McAlpine, G. et al., "Integrated Information Retrieval in a Knowledge Worker Support System", Proc. of the Intl. Conf. on Research and Development In Information Retrieval (SIGIR), Cambridge, MA, Jun. 25-28, 1989, Conf. 12, pp. 48-57.
- Tsuda, K. et al., "IconicBrowser: An Iconic Retrieval System for Object-Oriented Databases", Proc. of the IEEE Workshop on Visual Languages, Oct. 4, 1989, pp. 130-137.
- "Multiple Selection List Presentation Aids Complex Search", IBM Technical Disclosure Bulletin, vol. 36, No. 10, Oct. 1993, pp. 317-318.
- Han, J.: "Towards On-Line Analytical Mining in Large Databases" SIGMOD Record, Mar. 1998, ACM, USA, vol. 27, No. 1, pp. 97-107, XP000980233, ISSN: 0163-5808.

ART-UNIT: 2121

PRIMARY-EXAMINER: Starks, Jr.; Wilbert L.

ATTY-AGENT-FIRM: Bingham McCutchen LLP Marino; Fabio E.

ABSTRACT:

Some embodiments of the invention include methods for identifying clusters in a database, data warehouse or data mart. The identified clusters can be meaningfully understood by a list of the attributes and corresponding values for each of the clusters. Some embodiments of the invention include a method for scalable probabilistic clustering using a decision tree. Some embodiments of the invention, perform linearly in the size of the set of data and only require a single access to the set of data. Some embodiments of the invention produce interpretable clusters that can be described in terms of a set of attributes and attribute values for that set of attributes. In some embodiments, the cluster can be interpreted by reading the attribute values and attributes on the path from the root node of the decision tree to the node of the decision tree corresponding to the cluster. In some

embodiments, it is not necessary for there to be a domain specific distance function for the attributes. In some embodiments, a cluster is determined by identifying an attribute with the highest influence on the distribution of the other attributes. Each of the values assumed by the identified attribute corresponds to a cluster, and a node in the decision tree. In some embodiments, the CUBE operation is used to access the set of data a single time and the result is used to compute the influence and other calculations.

59 Claims, 13 Drawing figures

[Previous Doc](#)

[Next Doc](#)

[Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)[Search Forms](#)[Print Result Set](#)[Search Results](#)[Help](#)[User Searches](#)[Preferences](#)

DRAFT ENTRY 3 of 3

 [Generate Collection](#)[Print](#)[Logout](#)

File: USPT

May 13, 2003

DOCUMENT-IDENTIFIER: US 6564197 B2

TITLE: Method and apparatus for scalable probabilistic clustering using decision trees

Other Reference Publication (12):

Han, J.: "Towards On-Line Analytical Mining in Large Databases" SIGMOD Record, Mar. 1998, ACM, USA, vol. 27, No. 1, pp. 97-107, XP000980233, ISSN: 0163-5808.

[Previous Doc](#)[Next Doc](#)[Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)[Search Forms](#) [Generate Collection](#)[Print](#)[Search Results](#)[Help](#)[User Searches](#)

L15: Entry 21 of 25

File: USPT

Feb 13, 2001

[Preferences](#)

Logout NO: 6189004

DOCUMENT-IDENTIFIER: US 6189004 B1

TITLE: Method and apparatus for creating a datamart and for creating a query structure for the datamart

DATE-ISSUED: February 13, 2001

US-CL-CURRENT: 707/3; 707/102, 707/4

APPL-NO: 09/ 073753 [PALM]

DATE FILED: May 6, 1998

PARENT-CASE:

CROSS REFERENCES TO RELATED APPLICATIONS This application relates to the following group of applications. Each application in the group relates to, and incorporates by reference, each other application in the group. The invention of each application is assigned to the assignee of this invention. The group of applications includes the following. U.S. patent application Ser. No. 09/073,748, entitled "Method and Apparatus for Creating a Well-Formed Database System Using a Computer," filed May 6, 1998, and having inventors Craig David Weissman, Greg Vincent Walsh, and Eliot Leonard Wegbreit. U.S. patent application Ser. No. 09/073,752, entitled "Method and Apparatus for Creating and Populating a Datamart," filed May 6, 1998, and having inventors Craig David Weissman, Greg Vincent Walsh and Lynn Randolph Slater, Jr. U.S. patent application Ser. No. 09/073,733, entitled "Method and Apparatus for Creating Aggregates for Use in a Datamart," filed May 6, 1998, and having inventors Allon Rauer, Gregory Vincent Walsh, John P. McCaskey, Craig David Weissman and Jeremy A. Rassen. U.S. patent application Ser. No. 09/073,753, entitled "Method and Apparatus for Creating a Datamart and for Creating a Query Structure for the Datamart," filed May 6, 1998, and having inventors Jeremy A. Rassen, Emile Litvak, abhi a. shelat, John P. McCaskey and Allon Rauer.

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)[Search Forms](#) [Generate Collection](#) [Print](#)[Search Results](#)[Help](#)[User Searches](#)

L15: Entry 22 of 25

File: USPT

Dec 12, 2000

[Preferences](#)

Logout NO: 6161103

DOCUMENT-IDENTIFIER: US 6161103 A

TITLE: Method and apparatus for creating aggregates for use in a datamart

DATE-ISSUED: December 12, 2000

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Rauer; Allon	Mountain View	CA		
Walsh; Gregory Vincent	Cupertino	CA		
McCaskey; John P.	Mountain View	CA		
Weissman; Craig David	Belmont	CA		
Rassen; Jeremy A.	Sunnyvale	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Epiphany, Inc.	San Mateo	CA			02

APPL-NO: 09/ 073733 [PALM]

DATE FILED: May 6, 1998

PARENT-CASE:

CROSS REFERENCES TO RELATED APPLICATIONS This application relates to the following group of applications. Each application in the group relates to, and incorporates by reference, each other application in the group. The invention of each application is assigned to the assignee of this invention. The group of applications includes the following. U.S. patent application Ser. No. 09/073,748, entitled "Method and Apparatus for Creating a Well-Formed Database System Using a Computer," filed May 6, 1998, and having inventors Craig David Weissman, Greg Vincent Walsh and Eliot Leonard Wegbreit. U.S. patent application Ser. No. 09/073,752, entitled "Method and Apparatus for Creating and Populating a Datamart," filed May 6, 1998, and having inventors Craig David Weissman, Greg Vincent Walsh and Lynn Randolph Slater, Jr. U.S. patent application Ser. No. 09/073,733, entitled "Method and Apparatus for Creating Aggregates for Use in a Datamart," filed May 6, 1998, and having inventors Allon Rauer, Gregory Vincent Walsh, John P. McCaskey, Craig David Weissman and Jeremy A. Rassen. U.S. patent application Ser. No. 09/073,753, entitled "Method and Apparatus for Creating a Datamart and for Creating a Query Structure for the Datamart," filed May 6, 1998, and having inventors Jeremy A. Rassen, Emile Litvak, abhi a. shelat, John P. McCaskey and Allon Rauer.

INT-CL: [07] G06 F 17/30

US-CL-ISSUED: 707/4; 707/1, 707/3

US-CL-CURRENT: 707/4; 707/1, 707/3

FIELD-OF-SEARCH: 707/1-10, 707/200-208, 707/100-104

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected	Search ALL	Clear
---------------------------------	----------------------------	-----------------------

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>5386556</u>	January 1995	Hedin et al.	707/4
<u>5550971</u>	August 1996	Brunner et al.	707/3
<u>5659724</u>	August 1997	Borgida et al.	707/3
<u>5675785</u>	October 1997	Hall et al.	707/102
<u>5806060</u>	September 1998	Borgida et al.	707/3
<u>5995958</u>	November 1999	Xu	707/3

OTHER PUBLICATIONS

McAlpine, G. et al., "Integrated Information Retrieval in a Knowledge Worker Support System", Proc. of the Intl. Conf. on Research and Development in Information Retrieval (SIGIR), Cambridge, MA, Jun. 25-28, 1989, Conf. 12, pp. 48-57.

Tsuda, K. et al., "IconicBrowser: An Iconic Retrieval System for Object-Oriented Databases", Proc. of the IEEE Workshop on Visual Languages, Oct. 4, 1989, pp. 130-137.

"Multiple Selection List Presentation Aids Complex Search", IBM Technical Disclosure Bulletin, vol. 36, No. 10, Oct. 1993, pp. 317-318.

Kimball, R., "The Data Warehouse Toolkit", (1996) John-Wiley & Sons, Inc., 388 pages (includes CD ROM).

Chawathe, S. et al., "Change Detection in Hierarchically Structured Information", SIGMOD Record, vol. 25, No. 2, Jun. 1996, pp. 493-504.

Chawathe, S. et al., "Meaningful Change Detection in Structured Data", Proceedings of the 1997 ACM SIGMOD International Conference, ACM Press, 1997, pp. 26-37.

Labio, W. et al., "Efficient Snapshot Differential Algorithms for Data Warehousing", Department of Computer Science, Stanford University, (1996), pp. 1-13.

Wiener, J. et al., "A System Prototype for Warehouse View Maintenance", The Workshop on Materialized Views, pp. 26-33, Montreal, Canada, Jun. 1996.

Kawaguchi, A. et al., "Concurrency Control Theory for Deferred Materialized Views", Database Theory-ICDT '97, Proceedings of the 6th International Conference, Delphi, Greece, Jan. 1997, pp. 306-320.

Zhuge, Y. et al., "Consistency Algorithms for Multi-Source Warehouse View Maintenance", Distributed and Parallel Databases, vol. 6, pp. 7-40 (1998), Kluwer Academic Publishers.

Zhuge, Y. et al., "View Maintenance in a Warehousing Environment", SIGMOD Record, vol. 24, No. 2, Jun. 1995, pp. 316-327.

Wisdom, J. "Research Problems in Data Warehousing", Proc. of 4th Int'l Conference on Information and Knowledge Management (CIKM), Nov. 1995, 6 pages.

Yang, J. et al., "Maintaining Temporal Views Over Non-Historical Information Sources For Data Warehousing", Advances in Database Technology--EDBT '98, Proceedings of the 6th International Conference on Extending Database Technology, Valencia, Spain, Mar. 1998, pp. 389-403.

Quass, D., "Maintenance Expressions for Views with Aggregation", Proceedings of the 21st International Conference on Very Large Data Bases, IEEE, Zurich, Switzerland, (Sep. 1995), 9 pages.

Mumick, I. et al., "Maintenance of Data Cubes and Summary Tables in a Warehouse", Proceedings of the 1997 ACM SIGMOD International Conference, ACM Press, 1997, pp. 100-111.

Huyn, N., "Multiple-View Self-Maintenance in Data Warehousing Environments", Proceedings of the 23rd International Conference on Very Large Data Bases, IEEE, (1997), pp. 26-35.

Quass, D. et al., "Making Views Self-Maintainable for Data Warehousing", Proceedings of the Fourth International Conference, on Parallel and Distributes Information Systems, IEEE, Dec. 1996, pp. 158-169.

Gupta, H. "Selection of Views to Materialize in a Data Warehouse", Database Theory-ICDT '97, Proceedings of the 6th International Conference, Delphi, Greece, Jan. 1997, pp. 98-112.

Harinarayan, V. et al., "Implementing Data Cubes Efficiently", SIGMOD Record, vol. 25, No. 2, Jun. 1996, pp. 205-216.

Gupta, H. et al., "Index Selection for OLAP", IEEE Paper No. 1063-6382/97, IEEE (1997), pp. 208-219.

Labio, W. et al., "Physical Database Design for Data Warehouses", IEEE Paper No. 1063-6382/97, IEEE (1997), pp. 277-288.

Gupta, A. et al., "Aggregate-Query Processing in Data Warehousing Environments", Proceedings of the 21st VLDB Conference, Zurich, Switzerland, Sep. 1995, pp. 358-369.

O'Neill, P. et al., "Improved Query Performance with Variant Indexes", Proceedings of the 1997 ACM SIGMOD International Conference, ACM Press, 1997, pp. 38-49.

ART-UNIT: 271

PRIMARY-EXAMINER: Ho; Ruay Lian

ATTY-AGENT-FIRM: Wilson, Sonsini, Goodrich & Rosati

ABSTRACT:

A method for automatically defining aggregates for use in a datamart is described. The datamart includes fact and dimension tables. The method comprises accessing a schema description and an aggregates description for the datamart. The schema description specifies a schema, which in turn, defines the relationships between the fact tables and dimension tables of the datamart. The aggregates description specifies the aggregates, which define, from the schema definition, which aggregate tables are to be created from the fact tables and dimension tables in the datamart. The data in the aggregates correspond to the pre-computed results of specific types of queries. In response to a query, the aggregates can be searched to determine an appropriate aggregate to use in response to that query. The schema description is used to create a first set of commands to create and populate the fact and dimension tables. Additionally, a second set of commands to create, populate and access, the aggregates are also created from the aggregates description. Some of the commands of the first set of commands are executed causing the creation and population of the tables. Some of the commands of the second set of commands are executed causing the creation of the aggregate tables. Some of the remaining commands of the second set of commands are executed to populate the aggregate tables from the populated fact and dimension tables.

11 Claims, 43 Drawing figures

[Previous Doc](#)

[Next Doc](#)

[Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#) [Generate Collection](#) [Print](#)

L15: Entry 23 of 25

File: USPT

Dec 7, 1999

US-PAT-NO: 5999192

DOCUMENT-IDENTIFIER: US 5999192 A

TITLE: Interactive data exploration apparatus and methods

DATE-ISSUED: December 7, 1999

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Selfridge; Peter Gilman	Watchung	NJ		
Srivastava; Divesh	New Providence	NJ		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Lucent Technologies Inc.	Murray Hill	NJ			02

APPL-NO: 08/ 640411 [PALM]

DATE FILED: April 30, 1996

INT-CL: [06] G06 F 15/00

US-CL-ISSUED: 345/440

US-CL-CURRENT: 345/440

FIELD-OF-SEARCH: 395/140, 395/141, 395/142, 395/143, 345/440, 345/441, 345/433, 345/333, 345/334, 345/335

PRIOR-ART-DISCLOSED:

U. S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>5475851</u>	December 1995	Kodosky et al.	395/800
<u>5611059</u>	March 1997	Benton et al.	395/326
<u>5627979</u>	May 1997	Chang et al.	395/335

ART-UNIT: 272

PRIMARY-EXAMINER: Nguyen; Phu K.

ABSTRACT:

A data exploration tool which has a graphical user interface that employs directed graphs to provide histories of the data exploration operations. Nodes in the directed graphs represent operations on data; the edges represent relationships between the operations. One type of the directed graphs is the derivation graph, in which the root of the graph is a node representing a data set and an edge leading from a first node to a second node indicates that the operation represented by the second node is performed on the result of the operation represented by the first node. Operations include query, segmentation, aggregation, and data view operations. A user may edit the derivation graph and may select a node for execution. When that is done, all of the operations represented by the nodes between the root node and the selected node are performed as indicated in the graph. The operations are performed using techniques of lazy evaluation and encachement of results with the nodes. Another type of the directed graphs is the subsumption graph, in which an edge leading from a first node to a second node indicates that the second node stands in a subsumption relationship to the first node. If a result of the operation represented by the first node has been computed, the result is available to calculate the result of the operation represented by the second node.

32 Claims, 14 Drawing figures

[Previous Doc](#)[Next Doc](#)[Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)**End of Result Set**
 [Generate Collection](#) [Print](#)

L15: Entry 25 of 25

File: USPT

Jan 6, 1998

US-PAT-NO: 5706495

DOCUMENT-IDENTIFIER: US 5706495 A

**** See image for Certificate of Correction ****

TITLE: Encoded-vector indices for decision support and warehousing

DATE-ISSUED: January 6, 1998

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Chadha; Atul	Milpitas	CA		
Gupta; Ashish	Saratoga	CA		
Goel; Piyush	Monte Sereno	CA		
Harinarayan; Venkatesh	Stanford	CA		
Iyer; Balakrishna Raghavendra	San Jose	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE	CODE
International Business Machines Corporation	Armonk	NY			02	

APPL-NO: 08/ 643998 [PALM]

DATE FILED: May 7, 1996

INT-CL: [06] G06 F 17/30

US-CL-ISSUED: 395/602; 395/603

US-CL-CURRENT: 707/2; 707/3

FIELD-OF-SEARCH: 395/602, 395/601, 395/603

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<input type="checkbox"/> 5105353	April 1992	Charles et al.	395/700
<input type="checkbox"/> 5293616	March 1994	Flint	395/600
<input type="checkbox"/> 5495608	February 1996	Autoshenkov	395/603

5560007

September 1996

Thai

395/603

OTHER PUBLICATIONS

"Decision Support viewpoint: An Enterprise-wide Data Delivery Architecture," brochure, Microstrategy Incorporated, Vienna, VA, 1994, pp. 1-15.

"An Introduction to Multidimensional Database Technology," brochure, Kenan Technologies, Kenan Systems Corporation, Cambridge, MA, 1994, pp. 1-28.

"Red Brick High-Speed Query Accelerator of its Own," Computergram International, Dec. 15, 1994, (ISSN:0268-716x).

A. Shoshani. Statistical Databases: Characteristics, Problems and Some Solutions. Proceedings of the Eighth International Conference on Very Large Databases (VLDB), pp. 208-222, 1982.

RELEase 1.0, v91, n2, p1-27, Feb. 25, 1991 (ISSN:0740-935x).

Chang, W. Soliman, H.S. Sung, A.H., "Image Data Compression Using Counterpropagation Network," 1992 IEEE International Conference on Systems, Man and Cybernetics, (cat. No. 92CH3176-5) Oct. 18-21, 1992, pp. 405-409 vol. 1.

Frisch, Joseph, "Bit Vectors Vitalize Data Retrieval," Data Processing Magazine's Data Dynamics, vol. 13, No. 8 pp. 37-41, Aug./Sep. 1971.

Jacobsson, M., "Implementation of Comprssed Bit-Vector Indexes," Furo IFIP 79, North Holland Publishing Company, 1979, pp. 561-566.

Kimball, Ralph and Strehlo, Kevin, "Why Decision Support Fails and How to Fix it," Datamation, Jun. 1, 1994, pp. 40-45.

Marshall, Martin, "Data Warehouse Update to Include Bit-Mapped Indexing," CommunicationsWeek, No. 585, Nov. 20, 1995, p. 5.

Phillips, Ben, "Red Brick Props up Flagship Foundation," PC Week, vol. 12, No. 47, p. 45, Nov. 17, 1995.

"Multidimensional Analysis: Converting Corporation Data into Strategic Information," Arbor Software Corporation, Sunnyvale, CA.

J. Gray, A. Bosworth, A. Layaman and H. Pirahesh. "Data Cube: Relational Operator Generalizing Group-By, Cross-Tabs and Sub-Totals," IEEE, 1996, pp. 152-159.

Sybase's Fast Projection Index, "Faster Data Warehouses: New Tools Provide High Performance Querying through Advanced Indexing," InformationWeek Dec. 4, 1995, p. 77, ISSN: 8750-6874.

E.F. Codd, "Providing OLAP (On-line Analytical Processing) to User-Analysts: An IT Mandate," E.F. Codd and Associates, 1993.

"Decision Support Viewpoint: The Case for Relational OLAP," MicroStrategy, Inc, Vienna, Virginia, 1995, pp. 1-20.

ART-UNIT: 237

PRIMARY-EXAMINER: Kulik; Paul V.

ATTY-AGENT-FIRM: Merchant, Gould, Smith, Edell, Welter & Schmidt

ABSTRACT:

A method, apparatus, and article of manufacture for optimizing SQL queries in a relational database management system using a vectorized index. The vectorized index represents values in one or more of the columns of a particular table in the relational database. The vectorized index is comprised of a plurality of positions, wherein each of the positions comprises a linear array that represents a value for the specified columns in a corresponding row of the particular table in the relational database. To use the vectorized index, SQL operations are converted to a series of bit-vector operations on that index, where the result of the bit-vector operations is a list of row positions in the table.

72 Claims, 13 Drawing figures

[Previous Doc](#)

[Next Doc](#)

[Go to Doc#](#)